2

NUSC Technical Document 6902
1 June 1990

Recommendation Report for the Next-Generation Computer Resources (NGCR) Operating Systems Interface Standard Baseline

Operating Systems Standards Working Group (OSSWG)
Compiled by D. P. Juttelstad (NUSC)

Naval Underwater Systems Center
Newport, Rhode Island New London, Connecticut

PREFACE

This report was funded under NUSC Project No. A45146, "Next-Generation Computer Resources (NGCR)." The sponsoring activity is the Space and Naval Warfare Systems Command, through the work of the Operating Systems Standards Working Group (OSSWG). The OSSWG management structure is as follows:

NGCR Program Manager, H. Mendenhall (SPAWAR-324)

NGCR OSSWG Cochairman, CDR R. Barbour (SPAWAR-324)

NGCR OSSWG Cochairman, P. Oberndorf (NADC)

Approach Subgroup Chairman, T. Conrad (NUSC)

Requirements Subgroup Chairman, R. Bergman (NOSC)

Available Technology Subgroup Chairman, J. Oblinger (NUSC)

REVIEWED AND APPROVED: 1 JUNE 1990

P. A. La Brecque
Head, Combat Control Systems Department

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. AGENCY USE ONLY *(Leave blank)*

2. REPORT DATE
1 June 1990

3. REPORT TYPE AND DATES COVERED
Final

4. TITLE AND SUBTITLE
Recommendation Report for Next-Generation Computer Resources (NGCR) Operating Systems Interface Standard Baseline

5. FUNDING NUMBERS
PN A45146

6. AUTHOR(S)
Operating Systems Standards Working Group (OSSWG)

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Naval Underwater Systems Center
Newport Laboratory
Newport, RI 02841

8. PERFORMING ORGANIZATION REPORT NUMBER
TD 6902

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Space and Naval Warfare Systems Command
(SPAWAR-324)
Washington, DC 20363

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution is unlimited.

12b. DISTRIBUTION CODE

13. ABSTRACT *(Maximum 200 words)*

The Next-Generation Computer Resources (NGCR) Operating Systems Standards Working Group (OSSWG) conducted a survey of existing operating systems and operating systems interface standards to establish a baseline for the NGCR operating system interface. This report presents the results of that survey and the OSSWG recommendation for the standard baseline.

14. SUBJECT TERMS
Next-Generation Computer Resources
Operating Systems Interface

15. NUMBER OF PAGES
18

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	SAR

EXECUTIVE SUMMARY

The Next-Generation Computer Resources (NGCR) Operating Systems Standards Working Group (OSSWG) conducted a survey of existing operating systems and operating systems interface standards to establish a baseline for the NGCR operating system interface (OSIF). As a result of this survey, the total number of operating systems considered was reduced from 110 to 7, which then were formally evaluated. These seven were Alpha, ARTX, CRONUS, iRMX, Mach, ORKID, and POSIX.

The formal evaluation consisted of assessing the seven candidates against the requirements contained in the "NGCR OSSWG Requirements Document" (reference 1) and a set of eight programmatic issues. The numeric results of this evaluation identified three candidates as superior: Alpha, iRMX, and POSIX. To obtain a clear consensus of the OSSWG, an anonymous ballot was held that resulted in POSIX obtaining a 51-percent majority vote. Based on the results of the balloting, the NGCR OSSWG recommends POSIX be selected as the NGCR OSIF baseline. The working group also recommends that the Navy and OSSWG capitalize on the strengths of the other candidates, particularly Alpha and iRMX, in the continuing standards development.

TABLE OF CONTENTS

FOREWORD

The work reported herein was conducted over a period of a little more than a year by a joint team of Navy, other government, industry, and academic experts in the field of computer operating systems. Only a few of the Navy participants were actually funded to directly participate in this process. The superb accomplishments of the joint working group and its ability to complete the evaluation process in so short a time span derived from the dedication of all the participants to getting the job done. The outstanding contributions of all of the volunteers in this process are particularly noted and appreciated.

Special thanks are expressed to U.S. industry and academia for their staunch support of and participation in this working group. Their continued support and involvement are strongly solicited.

The OSSWG members who actively performed the evaluation of the final seven candidates were:

CDR Richard Barbour	SPAWAR-324
Richard Bergman	NOSC
Paul Bicknell	Mitre
Richard Brogan	Booz, Allen, & Hamilton
Dale Brouhard	NOSC
Gregory Bussiere	NUSC
Antonio Carangelo	Mitre
Gordon Caswell	ESL
Thomas Conrad	NUSC
B. Dasarathy	Concurrent Computer
Larry Daubert	Rockwell International
Isobel Davis	Raytheon
Steven Davis	DGM&S
Dr. Thomas Drake	Clemson University
Richard Dvorchak	Intel
LT Karl Fairbanks, Jr.	NWC
Gary Fisher	NIST
Lester Fraim	Honeywell
Dr. Karen Gordon	IDA
Dr. Mars Gralia	JHU/APL
Daniel Green	NSWC
Raymond Gretlein	Dynamics Research
Joseph Gwinn	Raytheon
Barbara Haleen	Unisys
James Hall	NIST
Neil Henderson	Litton Data Systems
Gail Holmes	NUSC
Steven Howell	NSWC
John Johnson	NAC
Daniel Juttelstad	NUSC
Kari Kruempel	Unisys
Dr. James Leathrum	Clemson University

Michael Linnig	Texas Instruments
Dr. Douglass Locke	IBM
Warren Loper	NOSC
Michael Morgan	PICHTR
Dr. John F. Nixon	General Electric Advanced Technology Laboratories
Patricia Oberndorf	NADC
James Oblinger	NUSC
Frank Prindle	NADC
John Reed	DEC
Carl Reinert	Computer Based Systems
Helmut Roth	NSWC
Dr. Timothy Saponas	Intel
John Shea	NOS
Del Swanson	Unisys
Maria Voreh	NRL
Patrick Watson	IBM

RECOMMENDATION REPORT FOR NEXT-GENERATION COMPUTER RESOURCES (NGCR) OPERATING SYSTEMS INTERFACE STANDARD BASELINE

1. INTRODUCTION

The charter of the Next-Generation Computer Resources (NGCR) Operating Systems Standards Working Group (OSSWG) is to establish a commercially acceptable operating system interface (OSIF) standard(s) for use in the development and deployment of Navy mission-critical computing systems in the mid-1990s and beyond. Candidates for NGCR standardization include existing public interface standards, as well as existing interface definitions (for example, based on commercial products or research prototypes) that could become public standards. The goal of the OSSWG is to do one of the following: (1) adopt existing standards/definition(s), if possible; (2) use Navy adaptations of existing standards/definitions; or (3) use Navy-created standards only if demanded by technical considerations (the worst case). Reference 1 contains the requirements of the baseline, while the goals and objectives of the NGCR program are cited in references 2 through 4. The specifics of these documents that pertain to the OSSWG are contained in the references 5 and 6.

This report summarizes the conclusions of the NGCR OSSWG and gives the recommendation for the OSIF baseline. Extensive details of the evaluation leading to this report are provided in references 7 and 8. A companion document, "After-Action Report for Next Generation Computer Resources Operating System Interface Baseline Selection Process" (reference 9), provides recommendations that resulted from the evaluation process.

Section 2 of the present report summarizes information about the 7 candidate operating systems selected from an initial field of 110. Section 3 identifies the final three candidates for selection as the OSIF baseline, along with their individual strengths, weaknesses, and risks. Section 4 provides the method by which the recommended OSIF baseline was selected and presents the OSSWG recommendation.

2. CANDIDATE SUMMARY

To establish a baseline for the NGCR OSIF, the OSSWG conducted a survey of existing operating systems and OSIF standards. As a result of the survey, the Available Technology (AT) Subgroup reduced the total number of candidates from 110 to 7, which then were evaluated by the entire OSSWG for selection as the OSIF baseline. More complete descriptions of the candidates can be found in reference 10. The seven candidates are described briefly in the following subsections.

2.1 ALPHA

The development of the Alpha operating system (OS) is currently being led by Concurrent Computer Corp., Westford, MA, but it is nonproprietary and in the public domain for U.S. Government use. Alpha is a distributed architecture kernel and object oriented. It manages all resources directly with application-specified actual time constraints, and includes real-time-distributed data management (e.g., transaction) mechanisms.

2.2 ARTX

The Ada Real-Time Executive (ARTX) is from Ready Systems, Palo Alto, CA. ARTX is an OS kernel, and it implements the full range of Ada semantic operations, including the complete Ada tasking model. Ready Systems also has RTAda-MP, which supports tightly-coupled multiprocessor systems (680 x 0 based).

2.3 CRONUS

CRONUS is a distributed OS being developed by Bolt Beranek and Newman (BBN), Cambridge, MA, and funded jointly by the Rome Air Development Center (RADC), the Naval Ocean Systems Center (NOSC), and the Air Force Electronic System Division (ESD). Designed to sit on top of heterogeneous local operating systems, CRONUS incorporates features such as heterogeneity, transparency, and object-oriented programming as well as more high-level features such as survivability, replication mechanism, multicluster and data base access, and distributed monitoring and control facilities.

To develop security mechanisms around CRONUS, the SDOS project is being pursued by Odyssey Research Associates, Ithaca, NY. This project also is funded by RADC.

2.4 iRMX

iRMX is an OS kernel developed commercially by Intel Corp., Beaverton, OR. It provides many standard kernel features in a mature and widely-used system. There are also other members of this family of operating systems, such as the Distributed iRMX, which provides distributed and transparent multiprocessing.

2.5 MACH

Mach is a multiprocessor-oriented OS kernel for a distributed environment being developed at Carnegie-Mellon University (CMU), Pittsburgh, PA. Funded by the Defense Advanced Research Projects Agency (DARPA), Mach approaches issues involved with multiprocessors, heterogeneity, transparency, and object-oriented programming.

The Trusted Mach project is another DARPA-sponsored research effort being pursued by Trusted Information Systems, Inc., Glenwood, MD. The goal of this project is to build a version of Mach that meets the B3 level of protection as specified in the "Department of Defense Trusted Computer System Evaluation Criteria" (reference 11). Current efforts have been concentrated at the kernel level, but other server levels have been defined for later efforts.

Mach is not presently a real-time system; however, CMU is in the process of developing real-time Mach. This development is being done by extending the existing Mach to include real-time features, such as real-time threads, an integrated real-time scheduler, and a real-time tool set.

2.6 ORKID

The Open Real-Time Kernel Interface Definition (ORKID) was developed by the VME International Trade Association (VITA), and has been represented at the NGCR OSSWG by Motorola. The objective of the ORKID standard is to provide a state-of-the-art, open, real-time kernel interface.

2.7 POSIX

The Portable Operating Systems Interface for Computer Environments (POSIX), a standards activity sponsored by the Institute of Electrical and Electronic Engineers (IEEE), is an attempt to define a standard OSIF and environment based on the UNIX operating system. There are several subgroups within IEEE Committee 1003 considering issues such as security, real-time verification, and Ada interfaces.

3. FINAL OSIF BASELINE CANDIDATES

3.1 SELECTION OF FINAL CANDIDATES

The seven candidate operating systems were reduced to the three candidates that consistently scored highest in the technical and programmatic criteria: Alpha, iRMX, and POSIX. Data analysis indicated that there was sufficient statistical significance to the gap between the scores of the top three and the scores of the other four candidates to justify reducing the eligible candidates for consideration to the top three. However, there was no sufficient statistical significance to the differences among the scores of the top three candidates to justify a selection among them. Therefore, the top three were evaluated with respect to:

- Strengths - Distinctive characteristics of the candidate that ameliorate its position for acceptance.

- Weaknesses - Deficienci s of the candidate that impose a known and predictable impact that is rectifiable.

- Risk - Deficiencies of the candidate that impose an unpredictable impact or result with respect to rectification; deficiencies whose rectification could destroy the underlying model.

Sections 3.2.1 through 3.2.3 discuss these areas in more detail. However, only the characteristics that distinguish one candidate from the other two are presented. Section 3.2.4 discusses general weaknesses and risks that apply to all three candidates.

3.2 EVALUATION CRITERIA FOR FINAL CANDIDATES

3.2.1 *Alpha*

3.2.1.1 ***Strengths.*** Alpha was designed with strict attention given to the areas of networks and communications, event and error management, resource management, and timing. The intent was to give users uniform and transparent network access to code, data, and objects within a system with a single, unified approach to managing these resources that meets real-time constraints. Resource management is done across all nodes within the system, also known as transnode resource management. Event and error management also is handled with a uniform approach with respect to all events within the system, including application events as well as time constraints. Users are given the ability to control and manage time by conveying to the kernel (through the interface) what their time constraints are. This directly affects the kernel's determination of the scheduling policy. These design goals for Alpha have been demonstrated and realized with prototypes by the Concurrent Computer Corp.

Because of another decision made during design, Alpha is strongly object oriented. This leads to extensibility, scalability, and an ordered design model for the OS. Alpha is extensible in that new objects can be created with specific attributes and added to the system as needed without change to any other existing objects. Also, other objects may communicate with the new object as well as with existing objects in the system. The scalability of Alpha refers to the ability to tailor Alpha by eliminating some unneeded objects without having detrimental effects on the operating system.

The inherent distributed orientation of Alpha along with its message-based threads and network transparency allow heterogeneous processors to coexist within the same system.

3.2.1.2 **Weaknesses.** The major weakness of Alpha is its small vendor and user base. Currently, there is only one announced vendor that is able to supply the Alpha system and another announced vendor that indicates it will have an Alpha interface in the future. There is a very small number of users actually using Alpha, although some companies have expressed interest in having an Alpha OS for their hardware. Also, no user working group has been formed for Alpha.

3.2.1.3 **Risks.** A major risk with selecting the Alpha kernel as a baseline OSIF is that there is no standards group or body formed to work on standardizing the Alpha interface. This presents a problem in that a standards organization would have to be found that would adopt the Alpha interface as its OSIF standard. Once this has been accomplished, there still is no assurance that the Alpha interface standard would be accepted by industry, possibly leaving the Navy in the same situation it faces today -- having to develop the OS itself.

There are also some concerns about Alpha's futuristic concept and its immaturity. An Alpha prototype exists, but it has not been implemented by many vendors on many different kinds of hardware. In addition, there are also few applications currently trying to make use of Alpha. Along these same lines is the question of whether the Alpha paradigm precludes efficient implementation. This has not been disproven.

Alpha, as submitted, does not represent a full local processor operating system (LPOS) application program interface (API). If Alpha were selected, there would be a large gap to fill within the Alpha interfaces to bring them up to a full LPOS API.

Alpha is based on a synchronous model. There was a concern, therefore, that if applications were to convert this model to an asynchronous one, the added overhead would be detrimental to real-time applications.

3.2.2 iRMX

3.2.2.1 Strengths. iRMX is the most mature of the three final candidates. This is apparent from the documentation that is user/reference oriented. These documentation follows a well-organized, consistent approach to describing iRMX. The presentation of the documentation follows the iRMX design-layered model.

The maturity of iRMX also is apparent in the user base. iRMX has more than 6000 real-time applications. With this user base, iRMX has demonstrated its maturity for providing a well-defined and understood conceptual model and has proven its interface capabilities to meet application needs.

iRMX is object oriented, providing an OS that is scalable and extensible and allowing for tailoring of the OS implementation for specific application needs. A user may create new object types and use system calls to manage these new objects.

The iRMX operating system is distributed: that is, the interface provides for message-passing communications between nodes. This capability has been proven through prototyping on a message-passing backplane, and the concept is believed to be extensible to local area networks as well.

3.2.2.2 Weaknesses. The only weakness unique in iRMX is that implementation information is interwoven with the interface description. (This is in reference to the RMX 386 kernel.) It was not apparent what the impact would be to the interface if the underlying processor was changed.

3.2.2.3 Risks. iRMX is a single-vendor product. It is not an accepted standard, nor is there a standards body considering iRMX as a potential standard interface. These factors present a major risk in getting the iRMX interface established as a commercially accepted standard. This operating system does have a substantial users group and a large user base, but it is specifically designed for Intel products. Additionally, there is the perspective that iRMX's technical similarity to other operating systems might make it unlikely as a standard the industry would accept.

Another risk is the lack of security. iRMX does not provide for security, and it is felt that the operational model may not be conducive to the addition of security.

3.2.3 POSIX

3.2.3.1 Strengths. The primary strengths of POSIX lie in its file management, synchronization, and scheduling interfaces. The POSIX file system consists of a full set of file functions that are consistent with the UNIX file system. Also, a full set of synchronization primitives exists in POSIX, and scheduling is extensible through the addition of policies.

POSIX is currently a standard and has an established user base because of its close relationship to UNIX. The standard is also commercially accepted, which makes POSIX very appealing to the Navy because many systems will be available for many types of hardware platforms. The standard has been proven to be implementable by many vendors, and most OS vendors have indicated their intent to, or currently do, supply a POSIX-compliant interface. Also, because POSIX is based on UNIX, it is a very mature technology. Limited risk exists because the interfaces have been proven to be implementable by many different system designers.

The POSIX interfaces are tailorable. Active profiling work within the standards body exists. This means that defining subsets and supersets of the standard is possible.

3.2.3.2 Weaknesses. The primary weaknesses of POSIX are in the areas of networking and communications, real-time capabilities, and distribution. Because networking and communications are not a principal part of the POSIX effort, a definite weakness in this area exists. Also, distribution is not a primary concern for the POSIX committee and, therefore, is a weak spot that must be addressed by the OSSWG. Real-time capability is sighted as a weakness because POSIX is intended to be a generic standard. Some real-time capabilities are precluded from the standard because of indications that the addition of these capabilities would dictate particular hardware support that may not be available to the OS designer. An example of this would be interrupt handling capabilities that are not included in the POSIX standard. At this time, there is no approved real-time standard within POSIX, although the 1003.4 subgroup is proceeding toward balloting and approval.

3.2.3.3 Risks. Three major risks are involved in choosing POSIX. The first is the question of how well the various subgroup standards (P1003.1 POSIX.1, P1003.4 Real-Time Extensions, P1003.5 Ada Bindings, and P1003.6 Security) will integrate. It is not guaranteed that these individual standards will work together and be able to be integrated into a single working standard.

The second risk with POSIX is how much influence the Navy can expect to have in the standards group. The Navy may not be able to persuade the POSIX committee to incorporate the changes essential to meet the Navy's needs. (An example of this issue is fault tolerance capabilities.)

Third, there was a concern that if applications were to convert this model to a synchronous one, the added overhead would be detrimental to real-time applications because POSIX is based on an asynchronous model.

3.2.4 *General Deficiencies*

3.2.4.1 General Weaknesses. Each of the top three candidates scored poorly in the service class of data interchange. This known deficiency in all three candidates does not appear to be a discriminating one; it is believed that extending all three candidates to meet the data interchange requirements will have an equal impact.

3.2.4.2 General Risks. The three final candidates have known deficiencies in the areas of:

- security;
- reliability adaptability and maintainability (also known as ***fault tolerance***); and
- Distribution performance.

Security is a risk for Alpha and iRMX because neither has addressed this issue. Security was addressed in the POSIX operating system, which scored well on the evaluation. However, security is considered a risk in POSIX, because it is not apparent that the P1003.6 (security subgroup) draft has been evaluated by the other working groups with respect to impact or feasibility. Of the three candidates, however, it is felt that the Alpha model is most compatible with security concepts.

None of the top three candidates properly addressed fault tolerance issues. The impact of extending each candidate to meet fault tolerance requirements is unknown.

Although implementation is not an issue of the OSIF, performance for real-time distributed systems is an objective. Based on the information provided by the candidates, it is felt that there is not enough information in the interfaces alone to have confidence in the ability of the underlying OS model to meet perceived Navy systems performance requirements.

None of the final three candidates meets the full set of Navy requirements as documented in reference 1, and each of the final candidates has various areas of strengths and weaknesses. Therefore, regardless of which candidate ultimately is selected, extensions will have to be made to the OSIF baseline before it meets the Navy's requirements. Extensions to the selected baseline should be approached by continuing to evaluate all the candidates to identify methods and approaches that might be appropriate for use in the extensions.

4. FINAL SELECTION PROCESS AND RECOMMENDATION

4.1 RATIONALE FOR SELECTING A SINGLE CANDIDATE

As discussed previously, the results of the technical evaluation by the OSSWG indicated that three candidates might be acceptable. A decision was made by the working group to validate the evaluation results of the final three candidates by examining these results carefully. This careful examination then would aid in interpreting these results. The scores on the eight programmatic issues also would be taken into account in this examination. If any of the three candidates proved to be acceptable from both a technical standpoint and from a programmatic standpoint, this candidate would be singled out by the working group for recommendation as the NGCR OSIF baseline. The decision to recommend only one candidate OS rather than multiple candidates as had originally been envisioned was made after much analysis and debate at two OSSWG meetings. In both meetings, the concept of pairing multiple candidates was rejected because the technical and programmatic results did not support such a pairing or composite solution. The OSSWG believed that, under the circumstances, a multiple candidate solution would be disadvantageous because this type of solution would (1) dilute NGCR resources and (2) fracture industry support. To maximize Navy influence and to achieve cost effectiveness, it was decided that the NGCR program should focus its efforts on a single candidate.

It should be noted that by recommending a single-candidate solution, the OSSWG is not ruling out the possibility of a family of standards. All that is ruled out is the possibility of a family based on an NGCR collection of divergent candidates. A family based on a single candidate is still possible and, in fact, probable.

At this point, the OSSWG believes that the family will take the form of a series or set of NGCR interface standards that can be tailored or scaled to any particular application (interface) requirements. It is anticipated that identifying the necessary extensions to the OSIF baseline selection to meet the overall requirements and defining the precise form of tailoring the OSIF baseline will be a major order of business for the next phase of the OSSWG.

4.2 FINAL CANDIDATE SELECTION

To make the selection of the NGCR OSIF baseline among the three finalists, it was anticipated that a definitive consensus could be obtained through an anonymous ballot at the NGCR OSSWG meeting held on 17-19 April 1990 at the Software Engineering Institute (SEI), Pittsburgh, PA. Prior to the voting, the detailed data analysis results and the strengths, weaknesses, and risks associated with the top three candidates were presented to the working group. An open discussion of this presentation was then held by the OSSWG, followed by the anonymous vote of the attending working group members.

There were two categories of voters at the meeting: eligible and advisory. Eligible voters were qualified members who had submitted their assessments during the evaluation process. In the absence of any eligible member, a representative of the absent person's organization was permitted to vote. The second category, or advisory voters, consisted of the remaining OSSWG members present who were allowed to submit votes.

A simple majority was established as grounds for a clear consensus from the qualified OSSWG members present. To distinguish eligible votes from advisory ones, each person present was given an envelope in which to place his ballot. Each person wrote his name and company on the envelope. If the individual was voting as a representative of a nonpresent eligible person, he also identified that individual on the envelope. The names on the envelopes were then sorted into eligible and advisory categories. The ballots were removed and counted to ensure the anonymity of the individual voters.

The voting results were as follows:

	Alpha	iRMX	POSIX	Total
Eligible	12	5	18	35
Advisory	3	1	6	10
Total	15	6	24	45

The POSIX operating system was given a 51-percent majority vote, making it the recommended candidate for selection as the NGCR OSIF baseline.

4.3 OSIF BASELINE RECOMMENDATION

Based on the results of the OSSWG evaluation process and a definitive consensus ballot, the NGCR OSSWG recommends that POSIX be selected as the NGCR OSIF baseline.

As a result of the evaluation, the OSSWG realized that none of the candidates totally meet the Navy's requirements for an OSIF standard. It was also apparent that individual candidates had varying strengths and weaknesses in the various requirement areas. With this in mind, the NGCR OSSWG further recommends that an approach be identified that maintains monitoring and evaluation of all the candidates, particularly Alpha and iRMX, so that in their evolution they may be used in expanding the capability of the selected baseline to meet the overall NGCR OSIF requirements.

REFERENCES

1. "NGCR OSSWG Requirements Document," SPAWAR-324 Informal Document, Version 2.0, 21 December 1989.

2. "Tentative Operational Requirements for Next-Generation Computer," CNO ltr to SPAWAR, Ser 098r5u353846, 31 December 1985, 31 December 1985.

3. "Operational Requirements for NGCR," CNO ltr to SPAWAR, Ser 098r/8u55086, 8 August 1988.

4. "NGCR Development Options Paper," SPAWAR ltr to NCGR OSSWG, Ser 324.253, 30 October 1987.

5. "White Paper on Network Operating Systems Standards," Naval Ocean Systems Center, San Diego, CA, August 1988.

6. "NGCR OSSWG Charter," SPAWAR-324 Informal Document, 16 March 1990.

7. "Evaluation Process Report for NGCR OSIF Baseline Selection." SPAWAR-324 Informal Document, Version 1.0, 7 May 1990.

8. "Evaluation Results Report for NGCR OSIF Baseline Selection," SPAWAR-324 Informal Document, Version 1.0, 7 May 1990.

9. J. T. Oblinger, Compiler, "After-Action Report for NGCR OSIF Selection Process," NUSC TD 6904, Naval Underwater Systems Center, Newport, RI, 1 June 1990.

10. "NGCR Available Technology Report," SPAWAR-324 Informal Document, Version 1.1, 18 April 1990.

11. "Department of Defense Trusted Computer System Evaluation Criteria," DoD 5200.28-STD, Department of Defense, Washington, DC, December 1985.

INITIAL DISTRIBUTION LIST

Addressee	No. of Copies
CNO (NOP-00, NOP-112, NOP-22, NOP-211, NOP-224, NOP-35, NOP-351, NOP-04, NOP-55, NOP-551, NOP-094, NOP-940, NOP-941, NOP-942, NOP-944, NOP-945, NOP-095, NOP-098, NOP-982, NOP-983)	20
CNR (OCNR-1133 (A. Van Tilborg))	2
U.S. Marine Corps HQ (C21, MCRDAC, PSE, LMA-1, CCA-51, OCIER, MCTSSA)	7
DIRSSP (NSP-00)	1
Cruise Missiles Project (Attn: C. Thomas, B. Tung)	2
NRL (Code 5540B (H. Lubbes), Code 5544 (M. Voreh))	2
COMNAVSECGRU	1
DONIRM (Attn: T. Senator)	1
PRESINSURV	1
INSURVLANT (Attn: Senior Member)	1
INSURVPAC (Attn: Senior Member)	1
NTISA, San Diego	1
COMNAVTELCOM	1
COMNAVAIRSYSCOM (NAIR-00 (30 copies), NAIR-5466 (B. Corson), NAIR-54664 (V. Skullman), PMA-209E (B. Anderson, M. Emerson))	34
COMSPAWARSYSCOM (SPAWAR-003-22, PD-40, PMW-141, PMW-142, PMW-143, PMW-144, PMW-145, PMW-146, PMW-147, PD-50, PMW-151, PMW-152, PMW-153, PMW-155, PMW-156, PMW-159, PD-60, PMW-161, PMW-162, PMW-163, PMW-164, PMW-610, PMW-620, PMW-630, PD-70, PD-70C, PD-70-1, PD-70-2, PD-70-3, PD-70-4, PMW-174, PD-80, PMW-180, PMW-181, PMW-182, PMW-183, PMW-184, PMW-162-13A (F. Deckelman), PMW-183-421 (R. Cockerill), SPAWAR-324 (P. Andrews, N. Stopyra), SPAWAR-324A (R. Barbour), SPAWAR-32431B (K. Chan), SPAWAR-3243 (A. Lewin, H. Mendenhall), SPAWAR-3243B (R. Owens, P. Pell), SPAWAR-3212 (M. Romeo), SPAWAR-32 (G. Wagner))	49
COMNAVSUPSYSCOM	1
COMNAVSEASYSCOM (SEA-00 (30 copies), SEA-90, SEA-06D1, PMS-412)	33
DTRC, Bethesda	3
DTRC, Annapolis	3
NADC (Code 7031 (P. Oberdorf), Code 7032 (F. Prindle), Code 7033 (C. Schmiedekamp), Library (3))	6
NCSC (Library (3))	3
NOSC (Code 424 (R. Bergman), Code 855 (W. Fitzgerald, R. Hall), Code 41 (A. Justice), Code 412 (W. Loper, D. Wilcox), Library (3))	9
NSWC, Dahlgren (Code N305 (D. Green), Library (3))	3
NSWC, White Oak (Code U33 (S. Howell, H. Roth), Library (3))	5
NAVWPNCEN (Code 3922 (K. Fairbanks), Code 3922 (C. Hall), Code 31C (J. Zenor), Library (3))	6
NSWSES	1
FCDSAA, San Diego (Code 8D (G. Robertson))	1
FCDSAA, Dam Neck	1
INTCOMBATSYSTESFAC	1
NAVELEXCEN, Charleston	1
NAVELEXCEN, Portsmouth	1

Addressee	No. of Copies
NAVELEXCEN, San Diego	1
NAVELEXCEN, Vallejo	1
NAVELEXCEN, Mayport	1
NAVELEXCEN, Pawtuxet River	1
NAVELEXCEN, Mechanicsburg	1
NAVELEXCENACT, St. Inigoes	1
NAVELEXCENACT, Philadelphia	1
NAVELEXSECCEN	1
NAVMASSO	1
NAVSPASYSACT	1
NATC (Code SY30 (J. Abrams, L. Campbell))	2
PMTC (Code 4003 (L. Cook))	1
NAVAVIONICCEN (Code 825 (J. Johnson))	1
COMNAVINTCOM	1
CNET	1
NETC	1
NETSAFA, Pensacola	1
NPS (Computer Science Dept, Administrative Science Dept, Engineering Dept, Library)	4
COMOPTEVFOR, Norfolk	1
ASSTSECNAV (S&L)	2
ASSTSECNAV (R,E,&S)	2
DSMC	1
APL, Johns Hopkins (Attn: B. Shaw, E. Conn, M. Gralia, G. Masson)	4
Commandant, U.S. Coast Guard (Code-G-FQA, G-TES (R. Buddeberg))	2
COMNAVDAC (Code 14 (P. Robinson))	2
NARDAC, Norfolk	1
NARDAC, Pensacola	1
NARDAC, San Diego	1
NARDAC, Alameda	1
NARDAC, Jacksonville	1
NARDAC, New Orleans	1
NARDAC, Pearl Harbor	1
NARDAC, Newport	1
NAVAIRENGCEN	1
NOS, Indian Head (Code 524)	1
NOS, Louisville (Code 50 (J. Shea))	1
DARPA (W. Scherlis)	1
USAF (AFSC/ENR) (Attn: P. Vaccaro)	1
USAF (SSC/XPT) (Attn: E. Crouse)	1
USAF (SAF/AQXA) (Attn: R. Gross)	1
Rome Air Development Center (COTD) (Attn: T. Lawrence)	1
USAISC (Attn: T. Fong)	1
NASA (Langley Research Center) (Attn: S. Beskenis)	1
Booz, Allen, & Hamilton (Attn: S. Alba, M. Karan, E. Rodriquez))	3
Institute for Defense Analysis (Attn: K. Gorden)	1
Arnold Associates (Attn: C. Arnold)	1
TRW Systems Division (Attn: S. Barry, A. Marmor-Squires)	2
Rockwell International (Attn: S. Bishop, L. Daubert, C. Dodd, F. Martin)	4

Addressee	No. of Copies
Carnegie-Mellon University (Attn: M. Borger, H. Tokuda)	2
Vitro (Attn: C. Brown, T. Darby, W. Greenspan, D. Sassaman)	4
CAE-Link Flight Simulation (Attn: L. Burns)	1
SOFTECH (Attn: B. Buseman)	1
MITRE (A. Carangelo)	1
IBM Manassas (Attn: D. Carrow, P. Watson)	2
ESL (Attn: G. Caswell)	1
UCLA (Attn: W. Chu)	1
DGM&S (Attn: S. Davis)	1
Telesoft (Attn: K. Dixon, R. Hadley)	2
G.E. Corporate Research & Development (Attn: W. Dixon)	1
Clemson University (Attn: T. Drake, J. Leathrum)	2
CEA (Attn: G. Driskill)	1
Intel (Attn: R. Dvorchak, T. Saponas)	2
Open Software Foundation (Attn: W. Finley)	1
Honeywell Federal System (Attn: L. Fraim)	1
DY-4 Systems (Attn: K. Clohessy, J. James)	2
Compumetrics (Attn: G. Coraluppi)	1
American Systems Corp. (Attn: M. Cornelison)	1
Martin Marietta Energy Systems (Attn: M. Cossey)	1
Control Data (Attn: R. Cunius, G. Jaeger)	2
Concurrent Computer (Attn: B. Dasarathy, E. Jensen)	2
Raytheon (Attn: I. Davis, J. Gwinn, L. Rainville)	3
Cray Research (Attn: C. Garrett)	1
Ready Systems (Attn: S. Gascon, J. Ready, D. Vrsalovic)	3
Computer Sciences (Attn: B. Gladstone, W. Lev, P. Wood)	3
Dynamics Research (Attn: R. Gretlein, G. Lampshire, P. Palatt)	3
Unisys (Attn: B. Haleen, M. Kamrad, K. Kruempel, J. Lonjers, H. Polzer, D. Swanson)	6
NIST Software Engineering (Attn: J. Hall, G. Fisher)	2
Integrated Software (Attn: S. Harbaugh)	1
Litton Data Systems (Attn: N. Henderson, A. Ruemke)	2
G.E. Ocean Systems Division (Attn: H. Hollander)	1
Texas Instruments (Attn: J. Jensen, K. Johnson)	2
UC Irvine (Attn: K. Kim)	1
Honeywell (Attn: D. Knudson)	1
Oracle (B. Kowalski, D. Reilly)	2
University of Pennsylvania (Attn: I. Lee, N. Prywes)	2
Electronic Data Systems (Attn: N. Lesher)	1
IBM (Attn: D. Locke)	1
Jet Propulsion Laboratory (Attn: F. Mathur, J. Rohr)	2
Advanced Technology & Research (Attn: R. Mayhall, A. Meskin)	2
SCTC Inc. (Attn: B. Miracle)	1
PICHTR (Attn: M. Morgan)	1
G.E. Advanced Technology Laboratories.(Attn: J. Nixon)	1
Numerix Federal Systems (Attn: D. Olson)	1
SYSCON (Attn: R. Owens)	1
Digital Equipment (Attn: J. Reed)	1
Computer-Based Systems (Attn: C. Reinert)	1
Hewlett-Packard (Attn: L. Reveron)	1

Addressee	No. of Copies
Teledyne Brown Engineering (Attn: E. Rozelle)	1
BBN Systems & Technologies (Attn: K. Schroder, S. Vinter)	2
Mnenomics (Attn: R. Semmes)	1
Industrial Programming (Attn: C. Sigda)	1
Logicon (Attn: B. Stauffer)	1
Trusted Information Systems (H. Tajalli)	1
Data General (Attn: L. Trumbore)	1
Motorola (Attn: R. Vanderlin)	1
IBM-FSD (Attn: D. Vogel)	1
Odyssey Research Associates (Attn: D. Weber)	1
Systems Exploration (Attn: B. Welty)	1
Westinghouse Electric (Attn: T. Wishart)	1
Advanced Systems Technologies (Attn: G. Wright)	1